Agentic AI in US Government: Optimization or Stranded Democratic Capital?

(The Democratic and Taxpayer Case for Agentic AI Removal)

Christopher Hunt Robertson, M.Ed.

(Written with the Support of Advanced AI - ChatGPT and Perplexity)

Earlier Title: "U.S. Democratic Restoration: Removing All Agentic AI from Most Domestic Government-Related Functions (All Levels - Except Military and Law Enforcement"

The introduction below has been added to Volume I of the book series in progress, *Non-Agentic Civic AI: Illuminating America's Civic Digital Field.*

A Respectful Note to the American Public

This book is written for you, with sincere hopes for your welfare and that of generations yet to come.

Not for machines. Not only for experts. Not only for those already inside government offices, technology firms, and policy institutions.

It is written for citizens of a constitutional republic who may increasingly find themselves governed, influenced, filtered, prioritized, scored, or administratively managed by systems they did not build, do not supervise, and may not fully understand.

In recent years, powerful institutions throughout the United States have begun integrating autonomous agentic artificial intelligence ("Agentic AI") into the machinery of government itself. These systems do not merely assist human beings in gathering information. Increasingly, they shape recommend-ations, prioritize responses, guide operational decisions, filter communications, identify risks, allocate attention, and influence how institutions understand the public they serve.

The promises surrounding these systems are familiar: greater efficiency, faster response, lower cost, smarter administration, predictive capability, optimized coordination, and more "effective" governance. Many intelligent and sincere people believe such systems will improve public life. Some applications may indeed appear useful for a time.

But this book begins from a different concern: history.

Again and again, free societies have embraced systems that appeared administratively irresistible - centralized technologies, emergency powers, surveillance architectures, bureaucratic shortcuts, and concentrated mechanisms of coordination - only to discover, often too late, that those systems gradually pushed democratic institutions beyond what they could sustainably bear.

The danger rarely arrived all at once.

It often accumulated quietly: through temporary exceptions, small delegations, efficiency incentives, institutional convenience, opaque procedures, and growing dependence upon systems that became progressively harder for ordinary citizens - and eventually even governments themselves - to meaningfully supervise.

Agentic AI in government possesses many of these same historical characteristics. It is easy to justify. Easy to scale. Easy to normalize. And increasingly difficult to see clearly once embedded deeply within institutional life.

By the time the public fully recognized what had changed, the erosion of accountability, visibility, and self-government was already well underway.

In a constitutional republic where political power regularly changes hands, this danger carries an additional and deeply destabilizing risk. Any system capable of shaping institutional judgment, prioritizing governmental attention, filtering civic visibility, or subtly influencing administrative outcomes will eventually be suspected of carrying the ideological assumptions, optimization priorities, or hidden biases of the previous administration that deployed it.

No lasting democratic trust can emerge from embedding systems that each successive government may reasonably fear were quietly tuned against it. Indeed, even systems perceived as "perfectly aligned" for one administration may become democratically destabilizing liabilities for the next.

A related and increasingly visible concern arises when public officials themselves are covertly influenced, improperly aligned, or quietly acting on behalf of foreign powers or concentrated private interests. In such cases, agentic systems wired into the machinery of government may gradually inherit and operationalize those hidden alignments at administrative scale, amplifying invisible human influence relationships while simultaneously making democratic responsibility harder for citizens to trace, understand, and publicly supervise. In addition, opaque proprietary systems, institutional dependency, and vendor-controlled optimization architectures can further limit meaningful public visibility, making it harder to distinguish ordinary commercial optimization from the gradual emergence of civic influence.

The people building and promoting these systems are not necessarily malicious. Many are attempting, in good faith, to solve real societal problems.

But good intentions do not preserve constitutional limits.

Institutional momentum does not guarantee democratic restraint.

And systems optimized primarily for speed, scale, prediction, and administrative efficiency do not naturally preserve the civic conditions upon which republican government ultimately depends: human judgment, visible accountability, public deliberation, constitutional restraint, and the moral responsibility carried by free citizens governing themselves. In a constitutional republic, some democratic friction is not failure. It is part of the infrastructure of legitimate self-government.

At present, the United States appears increasingly committed to a rushed deployment of Agentic AI. In other sectors, this may be helpful - but not in U.S. government. Agentic AI is already rapidly becoming part of the operational infrastructure of governance, administration, security, and public decision-making. This book does not deny that reality. Nor does it claim that agentic AI should be prevented from operating elsewhere in society.

It does insist, however, that Agentic AI must never become Civic AI operating near power. Systems capable of autonomous strategic action, adaptive optimization, operational drift, and increasingly opaque forms of machine-directed governance must not become the governing intelligence of the American republic itself.

Democratic restoration therefore requires removing Agentic AI from most government-related functions at all levels - except for our military and law enforcement agencies.

Civic legitimacy requires a different standard: refusal-based Non-Agentic Civic AI designed to illuminate public reality while remaining structurally incapable of partisan judgment, ideological steering, autonomous recommendation, or inherited political alignment. Here, refusal-based does not mean incapable. It means constitutionally constrained: designed to refuse delegated civic judgment, autonomous political authority, and the optimization of citizens themselves.

This is not a plan for smarter artificial rulers. Not a plan for replacing human conscience with optimized computation. Not a plan for quietly transferring democratic judgment into systems too complex, adaptive, opaque, or institutionally entrenched for ordinary citizens to meaningfully challenge.

Instead, the pages that follow propose a different kind of civic infrastructure: Non-Agentic Civic AI designed not to govern human beings, but to help human beings govern themselves more clearly. This alternative framework is visibility-first, not automation-first. Its purpose is not to replace human judgment, but to preserve and strengthen it. Its aim is to illuminate civic reality while keeping responsibility, legitimacy, and final authority fully human. Where democratic government succeeds, responsibility must remain human. Where democratic government fails, responsibility must remain human there as well.

A free people must never lose the ability to identify who governs, who decides, who bears responsibility, and who may be held accountable before the public.

You do not need to understand the technical mechanics of artificial intelligence to understand the central question this book asks: If Agentic AI in government is structurally prone to drift beyond meaningful human oversight, and if democratic responsibility cannot legitimately be delegated to artificial agents, are we wise to continue embedding such systems within the exercise of public authority without developing a more durable alternative?

This volume is my answer to that question.

It is an attempt to prepare, in advance, a workable democratic alternative before some future crisis forces the nation to confront these questions all at once.

Perhaps one morning, after a disturbing headline, institutional scandal, or visible failure of accountability, Americans will begin asking: How did we become so dependent upon systems we do not truly govern? Was there another path? What would democratic restoration require now? My hope is that this work may help answer those questions.

Throughout this work, you will encounter terms that may initially seem unfamiliar: "Non-Agentic Civic AI," "Civic Digital Visibility First," "The Fifty-Gates Democratic Permissions Threshold," and others.

These terms were not created merely as academic labels or proprietary concepts. They emerged gradually from the constitutional problems this work attempts to describe - problems for which our existing civic and technological vocabulary may not yet provide sufficiently precise public language. Sometimes new terms are needed when emerging technologies create distinctions older vocabularies cannot adequately express. Some readers may find these terms useful for public dialogue; others may prefer different language entirely. What matters is not whether these exact terms endure, but whether we find effective ways to approach the critical questions they are created to address.

For example, we must debate and decide: What are the constitutional boundaries between systems that assist human judgment and systems that increasingly substitute for it? Any successful transition to our Age of Humanity with AI will require an answer to this question. Free societies must preserve a meaningful boundary between systems that help citizens see and systems that begin quietly shaping, filtering, or replacing public judgment itself.

At the time of this writing, a growing number of institutions and researchers in fields such as constitutional law, AI safety, public administration, and cybersecurity have begun openly to confront a similar question: Where should the boundary lie between machine-assisted analysis and delegated machine judgment in public life?

Policy institutes, democracy advocates, technical governance bodies, public institutions, and AI developers have each, in their own language, begun describing the risks of allowing autonomous systems to assume roles once reserved for accountable human decision-makers.

No durable public consensus has yet emerged about how far such delegation should be permitted to go, or what constitutional limits a free society ought to place upon Agentic AI in government. *But we can all agree that democracies cannot govern what they cannot see.* Any serious democratic debate regarding boundaries must also confront a prior requirement: boundaries that cannot be clearly seen cannot be meaningfully compared, publicly judged, or constitutionally restrained. To deliberate responsibly about agentic boundaries, citizens, governments, institutions, and AI designers alike must first insist upon sufficiently shared visibility within our civic digital field. Before democratic alignment regarding such systems becomes possible, the boundaries themselves must become visible enough for a free people to understand, debate, and govern together.

The distinction between civic visibility infrastructure and delegated machine governance is not merely semantic. It is increasingly a structural question of democratic legitimacy.

Agentic AI in U.S. Government: Optimization or Stranded Democratic Capital?

One of the most practical democratic questions raised by the rushed deployment of Agentic AI in American government is not only whether such systems are innovative, but whether they are even durable. Even if many current officials, vendors, and reformers are acting with good intentions, public institutions may be investing billions of taxpayer dollars in architectures that prove legally unstable, constitutionally contestable, and institutionally difficult to sustain over time.

The modern case for Agentic AI in government is usually framed in the language of optimization, efficiency, modernization, and administrative speed. Yet democracies must ask a prior question before scaling such systems: Can a republic safely keep what it is being asked to buy? If a public technology later proves incompatible with due process, accountable human judgment, transparent administration, or constitutional limits on delegation, then the original expenditure may no longer appear as prudent modernization, but as a costly and avoidable misallocation of public funds.

This is where the danger of what may be called 'stranded democratic capital' becomes clear. The problem is not that government spent money on technology. The problem is that it may spend large public sums on systems that democracy itself cannot safely or legitimately keep - systems that later require suspension, litigation, redesign, or removal once their civic incompatibilities become more widely understood. Public officials possess a fiduciary duty not merely to pursue innovation, but to avoid binding taxpayers to governance architectures that a constitutional republic may ultimately prove unable to sustain. In that event, taxpayers may ultimately pay twice: first to install fragile Agentic AI infrastructures, and later to dismantle or replace them.

Recent public-sector studies already suggest reason for caution. Governments are rapidly expanding advanced Agentic AI use even while lagging behind on trustworthy governance foundations, transparent oversight structures, and durable accountability mechanisms. Other analysts have warned that agentic systems can encourage administrative rubber-stamping, diffuse responsibility across opaque algorithmic processes, and insert poorly understood forms of delegated power into institutions that remain constitutionally answerable to the public. If such concerns continue to deepen, many current investments may eventually come to resemble not durable civic infrastructure, but temporary experiments built upon democratically unstable ground.

The deeper issue is not technological capability alone, but whether democratic institutions can sustainably survive the forms of delegation such systems encourage. Throughout history, republics have repeatedly discovered that not every system that increases administrative efficiency is compatible with long-term civic legitimacy. Public institutions must therefore distinguish between technologies that strengthen democratic visibility and those that gradually displace accountable human judgment behind increasingly opaque systems of optimization.

A wiser course remains available. Democracies should invest in Non-Agentic Civic AI - civic visibility tools designed to illuminate conditions, strengthen shared public legibility, and support human judgment without replacing it. Such systems are far more likely to remain inspectable, lawful, publicly intelligible, and fiscally durable across generations. If democratic institutions are to modernize, they should do so in ways they can still justify, govern, audit, and afford decades later. In the end, the final measure will not be whether governments can build increasingly agentic systems. It is whether free societies can sustainably absorb them without undermining the constitutional trust, accountability, and human responsibility upon which democratic legitimacy ultimately depends.

As you consider the ideas herein, please remember:

Our democratic legacy allows us to question whether every capability that can be deployed in government should be deployed.

We are allowed to insist that constitutional self-government requires more than optimized administration.

We are allowed to demand systems that illuminate our civic field rather than quietly govern it.

We are allowed to require that where government acts, responsibility remains human and visible.

We are allowed to ask whether using taxpayer funds toward developing and installing Agentic AI under increasing democratic, legal, and political scrutiny is wise, given our current non-agentic alternatives.

And we are allowed, together, to restore democratic responsibility before it disappears beneath the logic of efficiency entirely.

The pages that follow are an invitation to do exactly that.

Our future need not be quietly decided for us by others.